

MAXIMUM CLIQUE and MINIMUM CLIQUE PARTITION in Visibility Graphs[★]

Stephan Eidenbenz and Christoph Stamm

Institute for Theoretical Computer Science, ETH Zürich, Switzerland
{eidenben, stamm}@inf.ethz.ch

Abstract In an alternative approach to “characterizing” the graph class of visibility graphs of simple polygons, we study the problem of finding a maximum clique in the visibility graph of a simple polygon with n vertices. We show that this problem is very hard, if the input polygons are allowed to contain holes: a gap-preserving reduction from the maximum clique problem on general graphs implies that no polynomial time algorithm can achieve an approximation ratio of $\frac{n^{1/8-\epsilon}}{4}$ for any $\epsilon > 0$, unless $NP = P$. To demonstrate that allowing holes in the input polygons makes a major difference, we propose an $O(n^3)$ algorithm for the maximum clique problem on visibility graphs for polygons without holes (other $O(n^3)$ algorithms for this problem are already known [3,6,7]). Our algorithm also finds the maximum weight clique, if the polygon vertices are weighted.

We then proceed to study the problem of partitioning the vertices of a visibility graph of a polygon into a minimum number of cliques. This problem is *APX*-hard for polygons without holes (i.e., there exists a constant $\gamma > 0$ such that no polynomial time algorithm can achieve an approximation ratio of $1 + \gamma$). We present an approximation algorithm for the problem that achieves a logarithmic approximation ratio by iteratively applying the algorithm for finding maximum weighted cliques. Finally, we show that the problem of partitioning the vertices of a visibility graph of a polygon *with* holes cannot be approximated with a ratio of $\frac{n^{1/14-\gamma}}{4}$ for any $\gamma > 0$ by proposing a gap-preserving reduction. Thus, the presence of holes in the input polygons makes this partitioning problem provably harder.

1 Introduction

Visibility problems have received considerable attention in the past. On the one hand, art gallery problems – such as MINIMUM VERTEX GUARD – have been studied intensively with respect to both, bounds on descriptonal complexity as well as computational complexity results. On the other hand, visibility graphs continue to draw interest. A *simple polygon with(out) holes* is given by its ordered sequence of vertices on the outer boundary, together with an ordered sequence

[★] We gratefully acknowledge the support of this work by the Swiss National Science Foundation.

of vertices for each hole, if any. Two polygon vertices *see* each other, iff the straight line segment connecting the two vertices does not intersect the exterior (or holes) of the polygon. A graph $G = (V, E)$ with vertices v_1, \dots, v_n is a *visibility graph*, iff there exists a simple polygon P (with or without holes) consisting of vertices p_1, \dots, p_n such that the polygon vertices p_i and p_j see each other, iff $(v_i, v_j) \in E$. The *visibility graph characterization* problem consists of finding a set of graph-theoretic properties that exactly define visibility graphs. It is closely related to the *visibility graph recognition* problem, which consists of determining if a given graph is a visibility graph. A lot of work has been done on the visibility graph characterization problem (see [15,13,24] or [25] for a survey), but it still is not satisfactorily solved. A different approach to “characterizing” the class of visibility graphs is to determine the computational complexity (and in case of NP -hardness the approximability) of classic graph-theoretic problems on visibility graphs. Actually, a considerable amount of work has been done that falls in the realm of this approach, because many classic graph-theoretic problems have a geometric interpretation in the context of visibility graphs. Also, the problem MINIMUM COLORING ON VISIBILITY GRAPH is mentioned as an open problem (with respect to its computational complexity) in an open problems list [22].

Consider, for example, the problem MAXIMUM INDEPENDENT SET ON VISIBILITY GRAPH, in which we are given a simple polygon with n vertices and we are to find the maximum independent set in the corresponding visibility graph. This problem corresponds to finding a maximum set of polygon vertices that are hidden from each other. The problem is therefore also called MAXIMUM HIDDEN VERTEX SET. It is known to be NP -hard [26], APX -hard for polygons without holes and hard to approximate with an approximation ratio of $\frac{n^{\frac{1}{6}} - \gamma}{4}$ for all $\gamma > 0$ for polygons with holes [12].

The problem MINIMUM DOMINATING SET ON VISIBILITY GRAPH corresponds to finding a minimum set C of polygon vertices such that each polygon vertex can be seen from at least one vertex in C . This problem is a variation of the well known art gallery problem MINIMUM VERTEX GUARD, which asks for a minimum number of vertices (guards) of a given polygon such that every point in the interior and on the boundary of the polygon can be seen from at least one guard. It is easy to see that the inapproximability results as well as approximability results for MINIMUM VERTEX GUARD carry over to MINIMUM DOMINATING SET ON VISIBILITY GRAPH, which therefore is APX -hard [11] for polygons without holes and not approximable with some approximation ratio that is logarithmic in the number of polygon vertices for polygons with holes [9]. Furthermore it is approximable with a logarithmic ratio [14].

In this paper we study the problem MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH(OUT) HOLES, in which we are given a simple polygon with(out) holes with n vertices and we are to find the largest clique in the corresponding visibility graph. We distinguish two separate problems by allowing holes or not. Note that in the case of polygons without holes, this problem corresponds to finding a largest (with respect to number of vertices) convex subpolygon of a given polygon. The geometric interpretation in the case of polygons with holes is unclear.

This problem has potential applications in the setting up of antenna networks in terrains (see [10,12] for the relationship of polygons with terrains), where all antennas must see each other in order to guarantee optimum connectivity.

We show that MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES cannot be approximated by any polynomial time algorithm with an approximation ratio of $\frac{n^{1/8-\epsilon}}{4}$ for any $\epsilon > 0$, unless $NP = P$ in Sect. 2. Thus, MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES is almost as hard to approximate as clique on general graphs. We propose a gap-preserving reduction (a technique introduced in [2]) from MAXIMUM CLIQUE on general graphs to get this result.

The problem MAXIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES is known to be solvable in time $O(n^3)$ by slightly adopting algorithms [3,6,7] that were developed to solve different problems (such as finding empty convex polygons that are maximum with respect to the number of vertices by connecting some of the input points). We propose an additional $O(n^3)$ algorithm for this problem for polygons without holes in Sect. 3, which uses dynamic programming. Our method also solves the weighted version of this problem, in which each vertex is assigned a weight value and the total weight of all vertices in the clique is to be maximized. We will use this weighted version (only with weights 0 and 1) to obtain an approximation algorithm for another problem (see Sect. 4).¹

This gap of “solvable in cubic time” vs. “almost as hard to approximate as clique” is the most extreme gap ever discovered between the two versions of a visibility problem on polygons with vs. without holes.

The problem MINIMUM CLIQUE PARTITION consists of finding a partitioning of the vertices of a given graph into a minimum number of disjoint vertex sets, each of which must be a clique in the graph. Again, we can define this problem on visibility graphs of polygons with or without holes. In the case of polygon without holes, this problem is closely related to the problem MINIMUM CONVEX COVER WITHOUT HOLES, which consists of covering a given polygon without holes with a minimum number of (possibly overlapping) convex polygons. MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPHS WITHOUT HOLES is a variant of MINIMUM CONVEX COVER WITHOUT HOLES, where only the vertices are of interest (not the edges or the interior area of the polygon).

A careful analysis (presented in [8]) of the reduction that was originally constructed to show the NP -hardness of MINIMUM CONVEX COVER [5] reveals that MINIMUM CONVEX COVER is APX -hard. The analysis can be easily adopted to work for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPHS WITHOUT HOLES. Therefore, MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPHS WITHOUT HOLES is APX -hard², i.e. there exists a constant $\epsilon > 0$ such that no polynomial time approximation algorithm can achieve an approximation ratio of $1+\epsilon$ for these problems, unless $NP = P$. In Sect. 4, we propose an approximation

¹ The fact that our $O(n^3)$ algorithm solves the *weighted* version of MINIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES, which will be used as a major building block for another approximation algorithm, is the main reason for including it in this paper, next to the obvious reason of self-containment.

² See [4] and [2] for an introduction to the class APX .

algorithm for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPHS WITHOUT HOLES that iteratively applies the algorithm for the weighted version of MAXIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES and show that it achieves a logarithmic approximation ratio. This result sheds some light on the approximability of MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPHS WITHOUT HOLES, but it still is not known whether a constant approximation ratio can be achieved or whether the logarithmic approximation algorithm presented is optimum.

There seems to be no straightforward geometric interpretation of MAXIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES, but the problem is certainly of theoretic interest, as we propose a gap-preserving reduction in Sect. 5 from MAXIMUM CLIQUE PARTITION on general graphs that shows that MAXIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES cannot be approximated with an approximation ratio of $\frac{n^{1/14-\gamma}}{4}$ for any $\gamma > 0$.

This is the first result for a visibility problem that is *NP*-hard no matter whether holes are allowed or not, where we are able to show that the approximation properties are clearly different for the cases of polygons with vs. without holes: While MAXIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES cannot be approximated with an approximation ratio of $\frac{n^{1/14-\gamma}}{4}$ for any $\gamma > 0$, we have a logarithmic approximation algorithm for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPHS WITHOUT HOLES.

In Sect. 6, we draw conclusions.

As for related work other than the previously mentioned, there are several surveys on art gallery and visibility problems [21] [25] [27]. As for computational complexity results, MINIMUM CONVEX COVER WITH(OUT) HOLES can be approximated with a logarithmic approximation ratio [8]. The problems MINIMUM VERTEX/EDGE/POINT GUARD, which are guarding problems with different types of guards, are known to be *NP*-hard [19] and *APX*-hard [11] for polygons without holes, and inapproximable with an approximation ratio logarithmic in the number of polygon vertices for polygons with holes [9]. Furthermore, MINIMUM VERTEX/EDGE GUARD can be approximated with a logarithmic approximation ratio for polygons with and without holes [14].

2 An Inapproximability Result for MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES

We propose a gap-preserving reduction from the MAXIMUM CLIQUE problem to the MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES problem. The technique of gap-preserving reductions [2] maps the promise problem of MAXIMUM CLIQUE to the promise problem MAXIMUM CLIQUE ON POLYGONS WITH HOLES. Suppose we are given an instance I of the promise problem MAXIMUM CLIQUE, i.e., a graph $G = (V, E)$ with $n := |V|$ and an integer k with $2 \leq k \leq n$, where $\epsilon > 0$ is arbitrarily small, but fixed. We are promised that the size of a maximum clique in the graph G is either at least k or strictly less than $\frac{k}{n^{1/2-\epsilon}}$. It




Figure 1. Basic construction: an input graph

is *NP*-hard to decide which of these two cases is true, because otherwise, MAXIMUM CLIQUE could be approximated by a polynomial time algorithm with an approximation ratio of $n^{1/2-\epsilon}$, which cannot be done unless $NP = P$ [16].

The basic idea of the reduction is shown in Figs. 1 and 2. For each instance I of MAXIMUM CLIQUE, i.e., for each graph $G = (V, E)$ with $n := |V|$ (as shown in an example in Fig. 1), we construct an instance I' of MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES, i.e., a polygon with holes (as shown in an example in Fig. 2). The main polygon is in the shape of a regular $2n$ -gon with vertices named v_i and v'_i for $i \in \{1, \dots, n\}$. For each vertex pair $(v_i, v_j) \notin E$, we construct two small triangular holes, one around the intersection point of the line segment from v_i to v_j and the line segment from v'_i to v'_{i+1} , and one around the intersection point of the line segment from v_i to v_j and the line segment from v'_j to v'_{j+1} . These triangular holes are designed to block the view of vertices v_i and v_j that are not supposed to see each other, since they are not connected by an edge in the input graph. The detailed, and rather technical construction of the holes is described in [12], and we therefore omit it here.

In order to make the reduction work, we refine the polygon with holes obtained thus far as follows:

For each vertex v_i let v_i^L (v_i^R) be the point on the line segment from v_i to v'_{i-1} (v'_i) that is closest to point v'_{i-1} (v'_i) such that the view of v_i^L to v_j^L (v_i^R to v_j^R) for all v_j is still blocked by the corresponding two holes, if vertices v_i and v_j are not connected in the input graph by an edge. These points are illustrated in Fig. 2.

For each vertex v_i , we replace the two line segments from v_i^L to v_i to v_i^R by a convex chain of $n^3 - 1$ line segments (called the *chain of* v_i). This is illustrated




Figure 2. Basic construction: polygon with holes resulting from the graph in Fig. 1

in Fig. 3. By the way that we chose points v_i^L and v_i^R , it is ensured that any two vertices from chains of v_i and v_j see each other, iff $(v_i, v_j) \in E$.

The following two lemmas allow us to prove the main result of this section. Let OPT denote the size of an optimum solution of the MAXIMUM CLIQUE instance I and let OPT' denote the size of an optimum solution of the MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES instance I' . Let $\epsilon > 0$.

Lemma 1. $OPT \geq k \implies OPT' \geq n^3 k$

Proof. If $OPT \geq k$, then there exists a clique of size k in I . We obtain a clique in I' of size $n^3 k$ by simply letting all the n^3 vertices of the chain of v_i be in the solution, if vertex $v_i \in V$ is in the clique. \square

Lemma 2. $OPT < \frac{k}{n^{1/2-\epsilon}} \implies OPT' < \frac{n^3 k}{n^{1/2-\epsilon}} + 3n^2$




Figure 3. Chain of vertex v_i

Proof. We prove the contraposition: $OPT' \geq \frac{n^3 k}{n^{1/2-\epsilon}} + 3n^2 \implies OPT \geq \frac{k}{n^{1/2-\epsilon}}$. Suppose we have a solution of I' with $\frac{n^3 k}{n^{1/2-\epsilon}} + 3n^2$ points. Since there are at most $n(n-1)$ holes with 3 vertices each and n additional vertices v'_i , there can be at most $3n(n-1) + n \leq 3n^2$ vertices in the clique that are not part of the chain of some v_i . Therefore, at least $\frac{n^3 k}{n^{1/2-\epsilon}}$ vertices of the clique must be in chains. Since a chain consists of only n^3 vertices, each chain can contribute at most n^3 vertices to the clique. Therefore, the number of chains that contain at least one point from the solution is at least $\frac{\frac{n^3 k}{n^{1/2-\epsilon}}}{n^3} = \frac{k}{n^{1/2-\epsilon}}$. Since no two vertices of two different chains v_i and v_j see each other unless $(v_i, v_j) \in E$, we immediately have a solution for I with at least $\frac{k}{n^{1-\epsilon}}$ vertices by letting v_i be in the clique if at least one point of the chain of v_i is in the solution. \square

Lemmas 1 and 2 transform the promise problem of MAXIMUM CLIQUE as mentioned above into a promise problem of MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES, where we are promised that an optimum solution contains either at least $n^3 k$ vertices or strictly less than $\frac{n^3 k}{n^{1/2-\epsilon}} + 3n^2$ vertices. It is also NP -hard to decide, which of the two cases is true, since otherwise, we could solve the NP -hard promise problem of MAXIMUM CLIQUE (see [2] for more details on the notion of such gap-preserving reductions). MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES can therefore not be approximated by any polynomial time approximation algorithm with an approximation ratio of:

$$\frac{n^3 k}{\frac{n^3 k}{n^{1/2-\epsilon}} + 3n^2} = \frac{n^3 k}{\frac{n^3 k + 3n^{1-2\epsilon}}{n^{1/2-\epsilon}}} \geq \frac{n^3 k}{\frac{2n^3 k}{n^{1/2-\epsilon}}} = \frac{n^{1/2-\epsilon}}{2}$$

We now need to express the size $|I'|$ of the MAXIMUM CLIQUE ON VISIBILITY GRAPH WITH HOLES instance I' by the size n of the MAXIMUM CLIQUE instance I . According to the construction, $|I'| \geq 2n^4$. We proceed:

$$\frac{n^{1/2-\epsilon}}{2} \geq \frac{|I'|^{\frac{1}{4}(\frac{1}{2}-\epsilon)}}{2^{\frac{1}{4}(\frac{1}{2}-\epsilon)}} \geq \frac{|I'|^{\frac{1}{8}-\frac{\epsilon}{4}}}{4}$$

This completes the proof of our main theorem of this section:

Theorem 1. MAXIMUM CLIQUE ON VISIBILITY GRAPHS WITH HOLES *cannot be approximated by any polynomial time algorithm with an approximation ratio of $\frac{|I'|^{1/8-\gamma}}{4}$, where $|I'|$ is the number of vertices in the polygon and where $\gamma > 0$, unless $NP = P$.*

3 An $O(n^3)$ Algorithm for MAXIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES

Our polynomial time algorithm for MAXIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES uses dynamic programming.

Suppose we are given a simple polygon P without holes, which consists of n vertices v_1, \dots, v_n in counterclockwise order. We first compute the visibility graph $G = G(P)$ of this polygon, which can be done in time $O(|E|)$, where E is the set of edges in G [17]. This allows us to answer queries of the form “Does vertex v_i see vertex v_j ?” in time $O(1)$. As we will use a weighted version of this problem to find an approximation algorithm for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITHOUT HOLES, we introduce a non-negative weight w_i for each vertex v_i . We are now to find a clique in G that has a maximum total weight. In the following, all operations are modulo n , where applicable. Let $A_{i,j,k}$ with $i < j \leq k$ be the *maximum* clique (with respect to its weight) among all cliques, which consist of vertices v_i, v_j and v_k and additional vertices $v_{j'}$ with $i < j' < j$. Let $|A_{i,j,k}|$ denote the weight of $A_{i,j,k}$. The optimum solution OPT is:

$$OPT = A_{i,j,j} \text{ where } i, j \text{ are such that } |A_{i,j,j}| = \max_{1 \leq i < j \leq n} |A_{i,j,j}|$$

Given all $A_{i,j,j}$, OPT can be computed in $O(n^2)$ time. A can be considered to be a three-dimensional table. It is initialized as follows:

$$A_{i,i+1,j} = \{v_i, v_{i+1}, v_j\}, \forall i, j, \text{ where vertices } v_i, v_{i+1}, v_j \text{ all see each other}$$

This initialization can be done in time $O(n^3)$. The remaining entries of the table A are initialized with empty sets and then computed according to Lemma 3.

Lemma 3. *Assume vertices v_i, v_j , and v_k see each other. Then, $A_{i,j,k} = A_{i,j',j} \cup v_k$, where j' is such that $|A_{i,j',j}| = \max |A_{i,j'',j}|$, where the maximum is taken over all j'' with $i \leq j'' \leq j$ and where $v_{j''}$ sees v_i, v_j , and v_k .*

Proof. The proof is inductive. Suppose we know that the lemma holds for $A_{i,j,k'}$ with $k' < k$. To show that it also holds for $A_{i,j,k}$, we assume by contradiction that there exists a clique P' , which consists of vertices v_i, v_j and v_k and additional vertices $v_{l'}$ with $i < l' < j$ and which is strictly heavier than $A_{i,j,k}$ (as computed in the Lemma).

Let v_l be the vertex in P' that is the neighbor of v_j in P' in clockwise order, when we interpret the clique P' as a convex polygon. Now, consider the clique




Figure 4. Proof of Lemma 3

$A_{i,l,j}$, which is maximum by assumption. Because v_j is the neighboring vertex of v_l in P' , we have $|P'| \leq |A_{i,l,j}| + w_k$. We will now argue that vertex v_k can be added to the clique $A_{i,l,j}$ and the resulting set of vertices (i.e. $A_{i,l,j} \cup v_k$) is still a clique.

Consider Fig. 4. First, note that vertex v_j must lie to the right of the line from v_i to v_k , because vertices v_i, v_j and v_k all see each other and because $i \leq j \leq k$. Since $v_i, v_l, v_j, v_k \in P'$ and $i \leq l \leq j \leq k$ and since P' is a clique, vertex v_l must lie to the right of the line from vertex v_i to v_k and to the left of the line from v_j to v_k . Now, consider all vertices $l'' \in A_{i,l,j}$ that lie between i and l (i.e. $i < l'' < l$). By definition of $A_{i,l,j}$, all these vertices see v_i, v_l and v_j . This implies that all vertices $v_{l''}$ also see v_k , because any polygon segment blocking the view of some vertex $v_{l''}$ to v_k would imply the existence of a polygon segment that would block the view of $v_{l''}$ to either v_i or v_l . We have shown that all vertices in $A_{i,l,j}$ also see v_k , therefore $A_{i,l,j} \cup v_k$ is a clique as well.

The polygon $A_{i,l,k}$ is among those polygons over which the maximum is taken in the Lemma to compute $A_{i,j,k}$. Therefore, $|A_{i,j,k}| \geq |P'|$, which is a contradiction to the assumption that P' is strictly heavier than $A_{i,j,k}$. \square

A trivial implementation of the algorithm thus suggested would have a running time of $O(n)$ for each of the $O(n^3)$ table entries, which results in an overall running time of $O(n^4)$. It is, however, possible to implement the algorithm with a total running time of $O(n^3)$. To achieve this, we show how to compute $A_{i,j,k}$ with i, j fixed and $A_{i,j',j}$ already computed for $i \leq j' \leq j$, in time $O(n)$ (for all k with $j \leq k \leq i$). This directly leads to an $O(n^3)$ algorithm, since there are only $O(n^2)$ pairs i, j .

To speed up the algorithm, fix i, j . Then compute all v_k with $j \leq k \leq i$ that are visible from v_i and v_j . Let K denote the counterclockwise ordered set of all these vertices v_k . Let L denote the clockwise ordered set of vertices v_l with $i \leq l \leq j$ that are visible to both v_i and v_j . For each vertex $v_l \in L$ (working from

v_j towards v_i): Determine, which vertices $v_k \in K$ are visible from v_l . Let $k' < k''$. Note that if v_l sees $v_{k'} \in K$, then it also sees $v_{k''} \in K$. Let $v_{k_{\min}}$ denote the first $v_k \in K$ that sees v_l . It suffices to just “link” $v_{k_{\min}} \in K$ to $A_{i,l,j}$ (depending on the implementation, a “link” could be an entry in some record field or a pointer). Note that as we work our way through L from v_j to v_i , the $v_{k_{\min}}$ ’s get smaller, i.e. proceed towards v_j . Thus, determining $v_{k_{\min}}$ can be done in total time $O(|K|)$ for all $v_l \in L$ (if $(|K| > |L|)$, otherwise it is $O(|L|)$). We now scan through K . If $v_k \in K$ is “linked” to some $A_{i,j,l}$, we compare the weight of $A_{i,j,l}$ with the weight of the currently optimum solution. If $|A_{i,j,l}|$ is greater than the weight of the currently optimum solution, we update the currently optimum solution to $A_{i,j,l}$. If v_k is not “linked”, we link it to the currently optimum solution. Now, set $A_{i,j,k}$ to the currently optimum solution with v_k added. We also store $|A_{i,j,k}|$. This scanning through K can be done in time $O(|K|)$. Thus, the total running time to compute $A_{i,j,k}$ for all k is $O(\max\{|L|, |K|\})$, which is $O(n)$.

Let us summarize the result of this section:

Theorem 2. *The weighted version of MAXIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES, where non-negative weights are assigned to the vertices, can be solved in time $O(n^3)$ using dynamic programming.*

4 An Approximation Algorithm for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITHOUT HOLES

Our approximation algorithm for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITHOUT HOLES iteratively applies the polynomial time algorithm for the weighted version of MAXIMUM CLIQUE ON VISIBILITY GRAPH WITHOUT HOLES. It works as follows for a given polygon P :

1. Compute the visibility graph $G(P)$ of the polygon P . Let all vertices have weight 1.
2. Find the maximum weighted clique C in $G(P)$ using the algorithm proposed in Sect. 3. Let all vertices $v_i \in C$ have weight 0. Add C to the solution S .
3. Repeat step 2 until there are no vertices with weight 1 left. Return S .

To obtain a performance guarantee of this algorithm, consider the MINIMUM SET COVER³ instance I , which has all polygon vertices v_i as elements and the vertices of each clique in the visibility graph of the polygon are a set in I . The greedy heuristic for MINIMUM SET COVER, which consists of recursively adding to the solution a set, which contains a maximum number of elements not yet covered by the solution, achieves an approximation ratio of $1 + \ln n$, where n is the number of elements in I [18]. Our algorithm works in exactly this way. Note that we do not have to compute all the sets of the MINIMUM SET COVER instance I (which would possibly be a number exponential in n), since it suffices to always

³ MINIMUM SET COVER consists of finding a minimum number of sets among a given collection of sets such that each element of a given universe appears in at least one of these sets.

compute a set (or clique), which contains a maximum number of vertices not yet covered by the solution, which is achieved by reducing the weights of the vertices already in the solution to 0. Thus, our algorithm is polynomial.

Theorem 3. MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITHOUT HOLES can be approximated with an approximation ratio of $1 + \ln n$, where n is the number of polygon vertices, by a greedy heuristic.

5 An Inapproximability Result for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES

MINIMUM CLIQUE PARTITION on general graphs is equivalent to MINIMUM GRAPH COLORING [4]. It cannot be approximated by any polynomial time algorithm with an approximation ratio of $n^{1/7-\epsilon}$, where $\epsilon > 0$ and n is the number of vertices in the graph [4]. We propose a gap-preserving reduction from MINIMUM CLIQUE PARTITION on general graphs to MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES.

Again, we map the *NP*-hard promise problem of MINIMUM CLIQUE PARTITION on general graphs, where we are promised that an optimum solution consists of either at most k or strictly more than $n^{1/7-\epsilon}k$ cliques, to a promise problem of MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES, where we are promised that an optimum solution consists of either at most $k + 3$ or strictly more than $n^{1/7-\epsilon}k$ cliques. We use the same construction as used in Sect. 2. However, we do not need to use the “chains” as introduced in Sect. 2. Let OPT (OPT') denote the size of an optimum solution of the MAXIMUM CLIQUE PARTITION (MAXIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES) instance I (I'). Let $\epsilon > 0$.

Lemma 4. $OPT \leq k \implies OPT' \leq k + 3$ and $OPT > n^{1/7-\epsilon}k \implies OPT' > n^{1/7-\epsilon}k$

Proof. For the first implication: If $OPT \leq k$, then there exists a solution of size k in I . We obtain a solution in I' of size $k + 3$ by simply letting all cliques from the solution in I be cliques in I' and by adding three more cliques. One of these consists of all the “bottom” vertices of all holes (i.e. those vertices that lie on line segments between points v'_{i-1} and v'_i for all i). The holes are constructed in such a way that these vertices actually form a clique (see [12]). The second clique consists of the “top” vertices of all holes. The third clique consists of all vertices v'_i . The construction of the reduction ensures that these additional cliques actually are cliques.

We prove the contraposition of the second implication: A solution for I' can be interpreted as a solution for I , where the additional vertices of I' are ignored. \square

We now proceed as in Sect. 2 using the same concepts. Lemma 4 and the fact that $|I'| \geq 3n^2$ allow us to prove:

Theorem 4. MAXIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITH HOLES cannot be approximated by any polynomial time algorithm with an approximation ratio of $\frac{|I'|^{1/14-\gamma}}{4}$, where $|I'|$ is the number of vertices in the polygon and where $\gamma > 0$, unless $NP = P$.

6 Conclusion

We have studied the two problems MAXIMUM CLIQUE ON VISIBILITY GRAPH and MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH for both polygons with and without holes. In the case of polygons without holes, the clique problem can be solved in polynomial time and this algorithm can be used in an approximation algorithm for the clique partition problem to achieve a logarithmic approximation ratio. The best inapproximability result known for the clique partition problem without holes is *APX*-hardness, thus the approximability of this problem is not yet precisely characterized.

In the case of polygons with holes, we have shown for both problems inapproximability ratios of n^ϵ for some $\epsilon > 0$, and have thus placed these two problems in the corresponding inapproximability class as defined in [2].

Our approach of “characterizing” the class of visibility graphs by studying classic graph problems for this class has been used before – at least implicitly. The computational complexity of the related problem of coloring the vertices of a visibility graph with a minimum number of colors is completely unknown and an open problem for future research [22]. Other open problems include, of course, determining the exact approximation threshold for MINIMUM CLIQUE PARTITION ON VISIBILITY GRAPH WITHOUT HOLES.

References

1. A. Aggarwal, S. Ghosh, and R. Shyamasundar; *Computational Complexity of Restricted Polygon Decompositions*; Computational Morphology, G. Toussaint, ed., North-Holland, pp. 1-11, 1988.
2. S. Arora and C. Lund; *Hardness of Approximations*; in: Approximation Algorithms for NP-Hard Problems (ed. Dorit Hochbaum), PWS Publishing Company, pp. 399-446, 1996.
3. D. Avis and D. Rappaport; *Computing the largest empty convex subset of a set of points*; Proc. 1st Ann. ACM Symposium Computational Geometry, pp. 161 - 167, 1985.
4. P. Crescenzi, V. Kann; *A Compendium of NP Optimization Problems*; in the book by G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, M. Protasi, *Complexity and Approximation. Combinatorial Optimization Problems and their Approximability Properties*, Springer-Verlag, Berlin, 1999; also available in an online-version at: <http://www.nada.kth.se/theory/compendium/>.
5. J. C. Culberson and R. A. Reckhow; *Covering Polygons is hard*; Proc. 29th Symposium on Foundations of Computer Science, 1988.
6. David P. Dobkin, Herbert Edelsbrunner, and Mark H. Overmars; *Searching for Empty Convex Polygons*; Algorithmica, 5, pp. 561 – 571, 1990.

7. Herbert Edelsbrunner, L. Guibas; *Topologically sweeping an arrangement*; J. Comput. System Sci. 38, pp. 165 – 194, 1989.
8. S. Eidenbenz and P. Widmayer; *An Approximation Algorithm for Minimum Convex Cover with Logarithmic Performance Guarantee*; manuscript, to appear.
9. S. Eidenbenz, C. Stamm, and P. Widmayer; *Inapproximability of some Art Gallery Problems*; Proc. 10th Canadian Conf. Computational Geometry, pp. 64-65, 1998.
10. S. Eidenbenz, C. Stamm, and P. Widmayer; *Positioning Guards at Fixed Height above a Terrain – An Optimum Inapproximability Result*; Lecture Notes in Computer Science, Vol. 1461 (ESA'98), pp. 187-198, 1998.
11. S. Eidenbenz; *Inapproximability Results for Guarding Polygons without Holes*; Lecture Notes in Computer Science, Vol. 1533 (ISAAC'98), pp. 427-436, 1998.
12. S. Eidenbenz; *How Many People Can Hide in a Terrain?*; Lecture Notes in Computer Science 1741 (ISAAC'99), pp. 184-194, 1999.
13. H. Everett, D. Corneil; *Negative Results on Characterizing Visibility Graphs*; Computational Geometry 5, pp. 51-63, 1995.
14. S. Ghosh; *Approximation algorithms for Art Gallery Problems*; Proc. of the Canadian Information Processing Society Congress, 1987.
15. S. Ghosh; *On Recognizing and Characterizing Visibility Graphs of Simple Polygons*; Discrete Comput Geom 17, pp. 143-162, 1997.
16. J. Hastad; *Clique is hard to approximate within $n^{1-\epsilon}$* ; Proc. of the Symposium on Foundations of Computer Science, 1996.
17. J. Hershberger; *Finding the Visibility Graph of a Polygon in Time Proportional to its Size*; Proc. of the 3rd Annual ACM Symposium on Computational Geometry, Waterloo, pp. 11 – 20, 1987.
18. D. Hochbaum; *Approximating Covering and Packing Problems: Set Cover, Vertex Cover, Independent Set, and Related Problems*; in: Approximation Algorithms for NP-Hard Problems (ed. Dorit Hochbaum), PWS Publishing Company, pp. 94-143, 1996.
19. D. T. Lee and A. K. Lin; *Computational Complexity of Art Gallery Problems*; IEEE Trans. Info. Th, pp. 276-282, IT-32, 1986.
20. J. O'Rourke and K. J. Supowit; *Some NP-hard Polygon Decomposition Problems*; IEEE Transactions on Information Theory, Vol IT-29, No. 2, 1983.
21. J. O'Rourke; *Art Gallery Theorems and Algorithms*; Oxford University Press, New York (1987).
22. J. O'Rourke; *Open problems in the combinatorics of visibility and illumination*; in Advances in Discrete and Computational Geometry, eds. B. Chazelle and J. E. Goodman and R. Pollack, (Contemporary Mathematics) American Mathematical Society, Providence, pp. 237-243, 1998.
23. C.H. Papadimitriou and M. Yannakakis; *Optimization, approximation, and complexity classes*; Proc. 20th ACM Symposium on the Theory of Computing, pp. 229 - 234, 1988.
24. L. Prasad, S.S. Iyengar; *A note on the combinatorial structure of the visibility graph in simple polygons*; Theoretical Computer Science 140, pp. 249-263, 1995.
25. T. Shermer; *Recent results in Art Galleries*; Proc. of the IEEE, 1992.
26. T. Shermer; *Hiding People in Polygons*; Computing 42, pp. 109-131, 1989.
27. J. Urrutia; *Art gallery and Illumination Problems*; in Handbook on Computational Geometry, edited by J.-R. Sack and J. Urrutia, 1998.